arXiv:1704.07816v1 [cs.CV] 25 Apr 2017

Introspective Classifier Learning: Empower Generatively

Long Jin Justin Lazarow Zhuowen Tu

Dept. of Computer Science and Engineering, Dept. of Cognitive Science
University of California, San Diego
$\{lojin, jlazarow, ztu\}$@$ucsd.edu$

Abstract

In this paper we propose introspective classifier learning (ICL) that emphasizes the importance of having a discriminative classifier empowered with generative capabilities. We develop a reclassification-by-synthesis algorithm to perform training using a formulation stemmed from the Bayes theory. Our classifier is able to iteratively: (1) synthesize pseudo-negative samples in the synthesis step; and (2) enhance itself by improving the classification in the reclassification step. The single classifier learned is at the same time generative — being able to directly synthesize new samples within its own discriminative model. We conduct experiments on standard benchmark datasets including MNIST, CIFAR, and SVHN using state-of-the-art CNN architectures, and observe improved classification results.

1. Introduction

In machine learning, great success has been achieved in obtaining powerful discriminative classifiers via supervised training. Other learning principles that are not fully supervised include unsupervised (Duda et al., 2000), semi-supervised (Zhu, 2005), weakly-supervised (Dietterich et al., 1997), and reinforcement (Sutton & Barto, 1998) which also point to promising directions when annotated data is limited.

Existing classifiers, such as decision trees (Quinlan, 1996), support vector machines (Vapnik, 1995), neural networks (LeCun et al., 1989), boosting (Freund & Schapire, 1997), and random forests (Breiman, 2001), carry out training processes to either implicitly or explicitly minimize an objective function that balances the training error and regularization. Recent studies reveal that even modern classifiers like deep convolutional neural networks (Krizhevsky et al., 2012) still make mistakes that look absurd to humans (Goodfellow et al., 2014b). A common way to improve the performance is by using more data, in particular "hard examples", to train the classifier. Different types of approaches have been proposed in the past including bootstrapping (Mooney et al., 1993), active learning (Settles, 2010), semi-supervised learning (Zhu, 2005), and data augmentation (Krizhevsky et al., 2012). However, the approaches above utilize data samples that are either already present in the given training set, or additionally created by humans or separate algorithms.

The key to granting the synthesis capability to a discriminative classifier is to make it internally generative. In the past, attempts have been made to build connections between generative models and discriminative classifiers (Friedman et al., 2001; Liang & Jordan, 2008; Jebara, 2012). Many existing approaches however combine discriminative classifiers with generative models to form hybrid models (Tu et al., 2008). In (Welling et al., 2002) a self supervised boosting algorithm was proposed to train a boosting algorithm by sequentially adding features as weak classifiers on additional self-generated negative samples; the generative discriminative modeling work in (Tu, 2007) generalizes the concept that a generative model can be successfully modeled by learning a sequence of discriminative classifiers via self-generated pseudo-negatives.

Inspired by the work that learns density functions (Welling et al., 2002) and generative models (Tu, 2007) using sequentially self-generated pseudo-negative samples, as well as recent success in deep learning (LeCun et al., 1989; Hinton et al., 2006; Krizhevsky et al., 2012; Gatys et al., 2015), we propose here an algorithm that capitalizes on the end-to-end learning of deep convolutional nets while making it internally generative. We specifically study how the generative aspect of our model can benefit its own discriminative training. We name our framework *introspective classifier learning* (ICL). There is a recent line of work to use a discriminative classifier to help with an external generator (Goodfellow et al., 2014a), which is different from our objective here. We aim at building a **single model** that is simultaneously discriminative and generative: being both a **discriminator** and a **generator** at the same time.

2. Significance and related work

The introspective classifier learning (ICL) framework being introduced here makes a number of contributions. (1) We introduce an end-to-end learning algorithm that is agnostic and both discriminative and generative under a single model (not a hybird one). (2) An efficient sampling procedure is developed to synthesize new data from a discriminative classifier. (3) A reclassification-by-synthesis algorithm is devised to iteratively augment negative samples and update the classifier. (4) We propose a formulation to seamlessly train a multi-class classifier on the given training set and augmented samples. We show consistent improvement over a state-of-the-art CNN classifier (ResNet (He et al., 2016a)) on CIFAR and SVHN datasets in the experiments.

To compare with the self-supervised boosting work (Welling et al., 2002), ICL has its advantage in: a). having a greater power for both discriminative and generative modeling beyond density estimation, b). being a direct discriminative classifier showing competitive experimental results of practical importance; c). being convolutional for automatic feature learning, and d). having a more efficient learning process. Compared with the generative model learning via discriminative classifier work (GDL) (Tu, 2007), ICL is advantageous for its: a). simplicity of having a single classifier as opposed to consisting of a sequence of boosting classifiers; b) being an effective discriminative classifier demonstrating direct improvement over state-of-the-art methods; c). greater representation power due to convolutional networks; and d). more efficient sampling process. Other methods (Gutmann & Hyvärinen, 2010; Tolstikhin et al., 2017) in which generative models are learned using discriminative classifiers share some common disadvantages as in (Welling et al., 2002; Tu, 2007).

Previous algorithms connecting generative modeling with discriminative classification (Friedman et al., 2001; Liang & Jordan, 2008; Tu et al., 2008; Jebara, 2012) fall in the category of hybrid models that are combinations of the two and mostly limited to specific tasks. Other directions in machine learning such as bootstrapping (Mooney et al., 1993), active learning (Settles, 2010; Beygelzimer et al., 2009), semi-supervised learning (Zhu, 2005), and data augmentation (Krizhevsky et al., 2012), try to effectively utilize the existing data whereas ICL is able to self-generate new data. Some existing work on introspective learning (Leake, 2012) has a different scope to the problem being tackled here.

Recent efforts in adversarial learning are also very interesting and worth mentioning. The work of (Goodfellow et al., 2014b) was motivated from an observation that adding small perturbations to an image leads to classification errors that are absurd to humans; their approach is however taken by augmenting positive samples from existing input whereas ICL is able to synthesize new samples from scratch. The generative adversarial nets algorithm (GAN) (Goodfellow et al., 2014a) engages two separate models, a generator and a discriminator, with the objective of making use of the discriminator to help the generator generate faithful samples; the discriminator in GAN is not meant to perform the generic two-class/multi-class classification task; thus, some special settings for semi-supervised learning (Goodfellow et al., 2014a; Radford et al., 2015; Zhao et al., 2016; Brock et al., 2016; Salimans et al., 2016) were created since the discriminators in GAN were trained to classify between "real" and "fake" samples. ICL instead has a single model that is both generative and discriminative, and thus, an improvement to ICL's generator leads to a direct means to ameliorate its discriminator; ICL also has a direct root in the Bayes theory from the derivation. Later development alongside GAN (Radford et al., 2015; Salimans et al., 2016; Zhao et al., 2016; Brock et al., 2016) share some similar aspects to GAN, which also do not achieve the same goal as ICL does. For example, an additional "not-real" class was generated in addition to the standard k classes in multi-class classification for a semi-supervised learning setting in (Salimans et al., 2016). ICL instead, maintains the same parameter setting in soft-max function identical to a standard multi-class CNN classifier.

Other generative modeling schemes such as the Mini-Max entropy theory (Zhu et al., 1997), inducing features (Della Pietra et al., 1997), auto-encoder (Baldi, 2012), Wake-sleep (Hinton et al., 1995), and recent CNN based generative modeling approaches (Xie et al., 2016b;a) are not for discriminative training and do not have a single model that is both generative and discriminative.

2.1. Relationship with GDL (Tu, 2007)

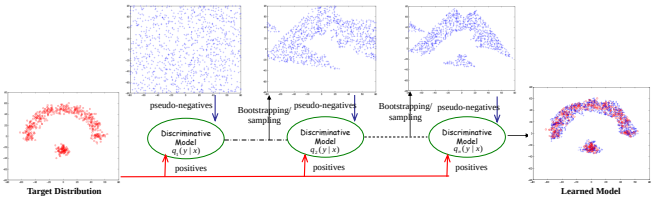

The pipeline of (Tu, 2007)

Figure 1. The pipeline of the GDL method developed in (Tu, 2007) (illustrated on a toy example) shows how to train a generative model via a sequence of discriminative classifiers using self-generated pseudo-negative samples.

The generative via discriminative learning framework (GDL) (Tu, 2007) learns a generative model through a sequence of discriminative classifiers (boosting (Freund & Schapire, 1997)) using repeatedly self-generated samples, called *pseudo-negatives*. Figure 1 shows the basic pipeline of GDL in (Tu, 2007). The far-left panel shows a set of

points (in red) as input data, and the task is to learn a generative model to characterize the distribution of these points, hence an unsupervised learning task; the GDL algorithm starts from a uniform distribution as a reference distribution to generate the first batch of pseudo-negatives (shown on the top-left panel) and a discriminative classifier is then trained to separate the given input data and the pseudo-negatives; new samples that pass the learned classifiers are considered as a new set of pseudo-negatives in the next round (shown on the top-middle panel) to train a new classifier; the algorithm repeats until the pseudo-negatives are no longer distinguishable from the input data (shown on the bottom-right panel); the learned generative model is then a concatenation of a series of discriminative classifiers learned through the process.

Our work is inspired by (Tu, 2007) but we also observe a number of areas for improvement to (Tu, 2007): (1) features in (Tu, 2007) are pre-specified manually instead of being automatically learned; (2) the sampling process is time-consuming, carried out by Markov chain Monte Carlo (MCMC) with a lengthy process to converge; (3) a series of discriminative classifiers are learned sequentially and those learned in the early stages do not get a chance to be updated later; (4) GDL was mainly for generative modeling and improvement to discriminative classification was not demonstrated in (Tu, 2007). In this paper, we pay our attention to the classification task and we focus on improving existing discriminative classification.

To summarize, we present a deep convolutional neural networks method that is simultaneously generative and discriminative under a single framework. We focus on the classification side and develop an effective algorithm for training. Encouraging experimental results have been attained on standard machine learning benchmark datasets.

3. Method

The pipeline of ICL is shown in Figure 2, which has an immediate improvement over GDL in several aspects that have been described in the previous section. One particular gain of ICL is its representation power and efficient sampling process through backpropagation as a variational sampling strategy.

3.1. Formulation

We start the discussion by introducing the basic formulation and borrow the notations from (Tu, 2007). Let $\mathbf{x}$ be a data sample (vector) and $y \in \{-1, +1\}$ be its label, indicating either a negative or a positive sample. In multi-class classification $y \in \{1, ..., K\}$. We study binary classification first. A discriminative classifier computes $p(y|\mathbf{x})$, the probability of $\mathbf{x}$ being positive or negative. $p(y = -1|\mathbf{x}) + p(y = +1|\mathbf{x}) = 1$. A generative model instead models $p(y, \mathbf{x}) = p(\mathbf{x}|y)p(y)$, which captures the underlying generation process of $\mathbf{x}$ for class y. In binary classification, positive samples are of primary interest. Under the Bayes rule:

$$p(\mathbf{x}|y = +1) = \frac{p(y = +1|\mathbf{x})p(y = -1)}{p(y = -1|\mathbf{x})p(y = +1)} p(\mathbf{x}|y = -1), \tag{1}$$

which can be further simplified when assuming equal priors $p(y = +1) = p(y = -1)$:

$$p(\mathbf{x}|y = +1) = \frac{p(y = +1|\mathbf{x})}{1 - p(y = +1|\mathbf{x})} p(\mathbf{x}|y = -1). \tag{2}$$

We make two interesting and important observations from eqn. (2): 1) $p(\mathbf{x}|y = +1)$ is dependent on the faithfulness of $p(\mathbf{x}|y = -1)$, and 2) a classifier C to report $p(y = +1|\mathbf{x})$ can be made **simultaneously generative and discriminative**. However, there is a requirement: having an informative distribution for the negatives $p(\mathbf{x}|y = -1)$ such that samples drawn $\mathbf{x} \sim p(\mathbf{x}|y = -1)$ have good coverage to the entire space of $\mathbf{x} \in \Re^m$, especially for samples that are close to the positives $\mathbf{x} \sim p(\mathbf{x}|y = +1)$, to allow the classifier to faithfully learn $p(y = +1|\mathbf{x})$. There seems to exist a dilemma. In supervised learning, we are only given a set of limited amount of training data and a classifier C is only focused on the decision boundary to separate the given samples and the classification on the unseen data may not be accurate. This can be seen from the top left plot in Figure 2. This motivates us to implement the synthesis part in learning: make a learned discriminative classifier generate samples that pass its own classification and see how different these generated samples are to the given positive samples. This allows us to attain a single model that has two aspects at the same time: a generative model for the positive samples and an improved classifier for the classification.

Suppose we are given a training set $S = \{(\mathbf{x}_i, y_i), i = 1..n\}$ and $\mathbf{x} \in \Re^m$ and $y \in \{-1, +1\}$. One can directly train a discriminative classifier C, e.g. a boosting algorithm (Freund & Schapire, 1997) or convolutional neural networks (LeCun et al., 1989) to learn $p(y = +1|\mathbf{x})$, which is always an approximation due to various reasons including insufficient training samples, generalization error, and classifier limitations. Previous attempts to improve classification by data augmentation were mostly done to increase the positive samples (Krizhevsky et al., 2012; Goodfellow et al., 2014b); we instead argue the importance of **increasing negative samples** to improve the classification performance. The dilemma is that $S = \{(\mathbf{x}_i, y_i), i = 1..n\}$ is limited to the given data. For clarity, we now use $p^-(\mathbf{x})$ to represent $p(\mathbf{x}|y = -1)$. Our goal is to gradually learn $p_t^-(\mathbf{x})$:

$$p_t^-(\mathbf{x}) \overset{t=\infty}{\rightarrow} p(\mathbf{x}|y = +1) \tag{3}$$

such that the samples drawn $\mathbf{x} \sim p_t^-(\mathbf{x})$ become indistinguishable from the given **positive** samples. We call the

samples drawn from $\mathbf{x} \sim p_t^-(\mathbf{x})$ **pseudo-negatives** (defined in (Tu, 2007)). We expand $S = \{(\mathbf{x}_i, y_i), i = 1..n\}$ by $S_e^t = S \cup S_{pn}^t$, where $S_{pn}^0 = \emptyset$ and for $t \geq 1$

$$S_{pn}^t = \{(\mathbf{x}_i, -1), i = n+1, ..., n+tl\}.$$

S_{pn}^t includes all the pseudo-negative samples self-generated from our model up to time t. l indicates the number of pseudo-negatives generated at each round. We define a reference distribution $p_r^-(\mathbf{x}) = U(\mathbf{x})$, where $U(\mathbf{x})$ is a Gaussian distribution (e.g. $\mathcal{N}(0.0, 0.3^2)$ independently). We carry out learning with $t = 0...T$ to iteratively obtain

$$q_t(y = +1|\mathbf{x}), \quad q_t(y = -1|\mathbf{x}) \tag{4}$$

by updating classifier C^t on $S_e^t = S \cup S_{pn}^t$. The initial classifier C^0 on $S_e^0 = S$ reports discriminative probability $q_0(y = +1|\mathbf{x})$. The reason for using q is because it is an approximation to the true p due to limited samples drawn in $\Re^m$. At each time t, we then compute

$$p_t^-(\mathbf{x}) = \frac{1}{Z_t} \frac{q_t(y = +1|\mathbf{x})}{q_t(y = -1|\mathbf{x})} p_r^-(\mathbf{x}), \tag{5}$$

where $Z_t = \int \frac{q_t(y=+1|\mathbf{x})}{q_t(y=-1|\mathbf{x})} p_r^-(\mathbf{x}) d\mathbf{x}$. Draw new samples

$$\mathbf{x}_i \sim p_t^-(\mathbf{x})$$

to expand the pseudo-negative set:

$$S_{pn}^{t+1} = S_{pn}^t \cup \{(\mathbf{x}_i, -1), i = n + tl + 1, ..., n + (t+1)l\}. \tag{6}$$

We name the specific training algorithm for our introspective classifier learning (ICL) framework **reclassification-by-synthesis**, which is described in Algorithm 1. We adopt convolutional neural networks (CNN) classifier to build an end-to-end learning framework with an efficient sampling process (to be discussed in the next section).

3.2. Reclassification-by-synthesis

We present our reclassification-by-synthesis algorithm for ICL in this section. A schematic illustration is shown in Figure 2. A single CNN classifier is being trained progressively, which is simultaneously a discriminator and a generator. With the pseudo-negatives being gradually generated, the classification boundary gets tightened, and hence improvement to the classifier performance. The reclassification-by-synthesis method is described in Algorithm 1. The key to the algorithm includes two steps: (1) reclassification-step, and (2) synthesis-step, which will be discussed in detail below.

3.2.1. Reclassification-step

The reclassification-step can be viewed as training a normal classifier on the training set $S_e^t = S \cup S_{pn}^t$ where $S = \{(\mathbf{x}_i, y_i), i = 1..n\}$ and $S_{pn}^0 = \emptyset$. $S_{pn}^t = \{(\mathbf{x}_i, -1), i = n+1, ..., n+tl\}$ for $t \geq 1$. We use CNN as our base classifier. When training a classifier C^t on S_e^t, we denote the parameters to be learned in C^t by a high-dimensional

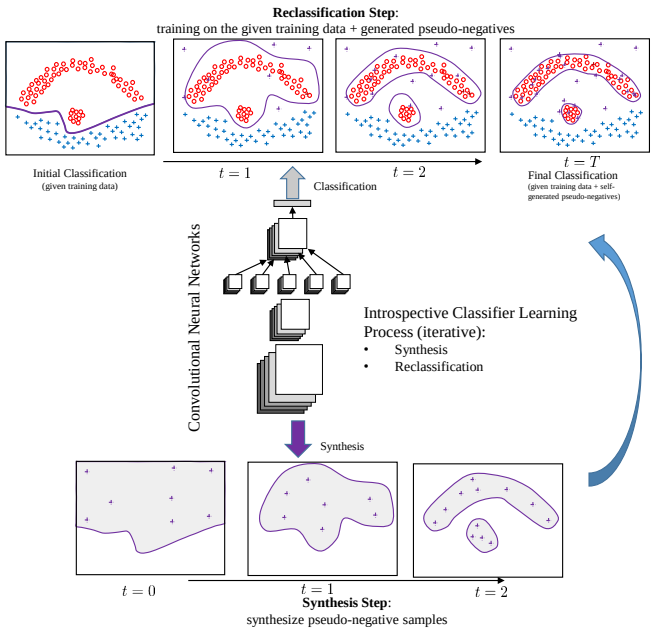

The reclassification-by-synthesis method for ICL training.

Figure 2. Schematic illustration of our proposed reclassification-by-synthesis algorithm for ICL training. The top-left figure shows the input training samples where the circles in red are positive samples and the crosses in blue are the negative samples. The bottom figures are the samples progressively self-generated by the classifier in the synthesis steps and the top figures show the decision boundaries (in purple) progressively updated in the reclassification steps. Pseudo-negatives (purple crosses) are gradually being generated and help tighten the decision boundary. We only maintain a single CNN classifier as opposed to having multiple boosting classifiers in GDL (see Figure 1).

vector $\mathsf{W}_t = (\mathbf{w}_t^{(0)}, \mathbf{w}_t^{(1)})$ which might consist of millions of parameters. $\mathbf{w}_t^{(1)}$ denotes the weights on the top layer combining the features $\phi(\mathbf{x}; \mathbf{w}_t^{(0)})$ and $\mathbf{w}_t^{(0)}$ carries all the internal representations. Without loss of generality, we assume a sigmoid function for the discriminative probability

$$q_t(y|\mathbf{x}; \mathsf{W}_t) = 1/(1 + \exp\{-y\mathbf{w}_t^{(1)T} \phi(\mathbf{x}; \mathbf{w}_t^{(0)})\}),$$

where $\phi(\mathbf{x}; \mathbf{w}_t^{(0)})$ defines the feature function for $\mathbf{x}$. Both $\mathbf{w}_t^{(1)}$ and $\mathbf{w}_t^{(0)}$ can be learned by the standard stochastic gradient descent algorithm via backpropagation to minimize a cross-entropy loss with an additional term on the pseudo-negatives:

$$\mathcal{L}(\mathsf{W}_t) = - \sum_{(\mathbf{x}_i, y_i) \in S}^{i=1..n} \ln q_t(y_i|\mathbf{x}_i; \mathsf{W}_t) - \sum_{(\mathbf{x}_i, -1) \in S_{pn}^t}^{i=n+1..n+tl} \ln q_t(-1|\mathbf{x}_i; \mathsf{W}_t). \tag{7}$$

Algorithm 1 Outline of the reclassification-by-synthesis algorithm for discriminative classifier training.

Input: Given a set of training data $S = \{(\mathbf{x}_i, y_i), i = 1..n\}$ with $\mathbf{x} \in \Re^m$ and $y \in \{-1, +1\}$.
Initialization: obtain a distribution the negative samples: $p_r^-(\mathbf{x}) = U(\mathbf{x})$ and train an initial CNN binary classifier C^0 on S, $q_0(y = +1|\mathbf{x})$. $S_{pn}^0 = \emptyset$. $U(\mathbf{x})$ is a zero mean Gaussian distribution.
For t=0..T
1. Update the model: $p_t^-(\mathbf{x}) = \frac{1}{Z_t}\frac{q_t(y=+1|\mathbf{x})}{q_t(y=-1|\mathbf{x})}p_r^-(\mathbf{x})$.
2. Synthesis-step: sample l pseudo-negative samples $\mathbf{x}_i \sim p_t^-(\mathbf{x}), i = n+tl+1, ..., n+(t+1)l$ from the current model $p_t^-(\mathbf{x})$ using a variational sampling procedure (stochastic gradient on the input).
3. Augment the pseudo-negative set with $S_{pn}^{t+1} = S_{pn}^t \cup \{(\mathbf{x}_i, -1), i = n+tl+1, ..., n+(t+1)l\}$.
4. Reclassification-step: Update CNN classifier to C^{t+1} on $S_e^{t+1} = S \cup S_{pn}^{t+1}$, resulting in $q_{t+1}(y = +1|\mathbf{x})$.
5. $t \leftarrow t+1$ and go back to step 1 until convergence (e.g. no improvement on the validation set).
End

3.2.2. SYNTHESIS-STEP

In the reclassification step, we obtain $q_t(y|\mathbf{x}; \mathsf{W}_t)$ which is then used to update $p_t^-(\mathbf{x})$ according to eqn. (5):

$$p_t^-(\mathbf{x}) = \frac{1}{Z_t}\frac{q_t(y=+1|\mathbf{x}; \mathsf{W}_t)}{q_t(y=-1|\mathbf{x}; \mathsf{W}_t)}p_r^-(\mathbf{x}). \tag{8}$$

In the synthesis-step, our goal is to draw fair samples from $p_t^-(\mathbf{x})$. In (Tu, 2007), various Markov chain Monte Carlo techniques (Liu, 2008) including Gibbs sampling and Iterated Conditional Modes (ICM) have been adopted, which are often slow. Motivated by the DeepDream code (Mordvintsev et al., 2015) and Neural Artistic Style work (Gatys et al., 2015), we update a random sample $\mathbf{x}$ drawn from $p_r^-(\mathbf{x})$ by increasing $\frac{q_t(y=+1|\mathbf{x}; \mathsf{W}_t)}{q_t(y=-1|\mathbf{x}; \mathsf{W}_t)}$ using backpropagation. Note that the partition function (normalization) Z_t is a constant that is not dependent on the sample $\mathbf{x}$. Let

$$g_t(\mathbf{x}) = \frac{q_t(y=+1|\mathbf{x}; \mathsf{W}_t)}{q_t(y=-1|\mathbf{x}; \mathsf{W}_t)} = \exp\{\mathbf{w}_t^{(1)T}\phi(\mathbf{x}; \mathbf{w}_t^{(0)})\}, \tag{9}$$

and take its ln, which is nicely turned into the logit of $q_t(y = +1|\mathbf{x}; \mathsf{W}_t)$

$$\ln g_t(\mathbf{x}) = \mathbf{w}_t^{(1)T} \cdot \phi(\mathbf{x}; \mathbf{w}_t^{(0)}). \tag{10}$$

Starting from $\mathbf{x}$ drawn from $p_r^-(\mathbf{x})$, we directly increase $\mathbf{w}_t^{(1)T}\phi(\mathbf{x}; \mathbf{w}_t^{(0)})$ using stochastic gradient ascent on $\mathbf{x}$ via backpropagation, which allows us to obtain fair samples subject to eqn. (8). Gaussian noise can be added to eqn. (10) along the line of stochastic gradient Langevin dynamics (Welling & Teh, 2011).

Sampling strategies When conducting experiments, we carried out several strategies using stochastic gradient descent algorithm (SGD) including: i) early stopping for the sampling process after $\mathbf{x}$ becomes positive (or a fixed small number); ii) sampling for a fixed, large number of steps. We found early-stopping to achieve a good balance between effectiveness and efficiency, which is aligned with contrastive divergence (Carreira-Perpinan & Hinton, 2005) where a short Markov chain is simulated.

Building the connection between SGD and MCMC is an active area in machine learning (Welling & Teh, 2011; Chen et al., 2014; Mandt et al., 2017). In (Welling & Teh, 2011), combining SGD and additional Gaussian noise under annealed stepsize results in a simulation to Langevin dynamics MCMC. SGD is also related to the diffusion process in MCMC (Grenander & Miller, 1994). A recent work (Mandt et al., 2017) further shows the similarity between constant SGD and MCMC, along with analysis of SGD using momentum updates. Our progressively learned discriminative classifier can be viewed as carving out the feature space on $\phi(\mathbf{x})$, which essentially becomes a equivalent class for the positives; the volume of the equivalent class that satisfies the condition is exponentially large, as analyzed in (Wu et al., 2000). The probability landscape of positives (equivalent class) makes our SGD sampling process not particularly biased towards a small limited modes. Results in Figure 3 illustrates that large variation of the sampled/synthesized examples.

3.3. Multi-class classification

One-vs-all

In the above section, we discussed the binary classification case. When dealing with multi-class classification problems, such as MNIST and CIFAR-10, we will need to adapt our proposed reclassification-by-synthesis scheme to the multi-class case. This can be done directly using a one-vs-all strategy by training a binary classifier C_i using the $i-th$ class as the positive class and then combine the rest classes into the negative class, resulting in a total of K binary classifiers. The training procedure then becomes identical to the binary classification case. If we have K classes, then the algorithm will train K individual binary classifiers with

$$< (\mathbf{w}_t^{(0)_1}, \mathbf{w}_t^{(1)_1}), ..., (\mathbf{w}_t^{(0)_K}, \mathbf{w}_t^{(1)_K}) > .$$

The prediction function is simply

$$f(\mathbf{x}) = \arg\max_k \exp\{\mathbf{w}_t^{(1)_k T} \cdot \phi(\mathbf{x}; \mathbf{w}_t^{(0)_k})\}.$$

The advantage of using the one-vs-all strategy is that the algorithm can be made nearly identical to the binary case at the price of training K different neural networks.

Softmax function

It is also desirable to build a single CNN classifier to perform multi-class classification directly. Here we propose a new formulation to train an end-to-end multiclass classifier

directly. Since we are directly dealing with K classes, the pseudo-negative data set will be slightly different and we introduce negatives for each individual class by $S^0_{pn} = \emptyset$ and:

$$\begin{aligned} S^t_{pn} &= \{(\mathbf{x}_i, -k), k = 1, ..., K, \\ &\quad i = n + (t-1) \times k \times l + 1, ..., n + t \times k \times l\} \end{aligned}$$

Suppose we are given a training set $S = \{(\mathbf{x}_i, y_i), i = 1..n\}$ and $\mathbf{x} \in \Re^m$ and $y \in \{1, ..K\}$. We want to train a single CNN classifier with

$$W_t =< \mathbf{w}^{(0)}_t, \mathbf{w}^{(1)_1}_t, ..., \mathbf{w}^{(1)_K}_t >$$

where $\mathbf{w}^{(0)}_t$ denotes the internal feature and parameters for the single CNN, and $\mathbf{w}^{(1)_k}_t$ denotes the top-layer weights for the $k-th$ class. We therefore minimize an integrated objective function

$$\begin{aligned} \mathcal{L}(\mathsf{W}_t) = &- \sum_{i=1}^{n} \ln \frac{\exp\{\mathbf{w}^{(1)_{y_i}T}_t \cdot \phi(\mathbf{x}_i; \mathbf{w}^{(0)}_t)\}}{\sum_{k=1}^{K} \exp\{\mathbf{w}^{(1)_k T}_t \cdot \phi(\mathbf{x}_i; \mathbf{w}^{(0)}_t)\}} \\ &+ \sum_{i=n+1}^{n+t\times K\times l} \ln(1 + \exp\{\mathbf{w}^{(1)_{|y_i|}T}_t \cdot \phi(\mathbf{x}_i; \mathbf{w}^{0}_t)\}) \quad (11) \end{aligned}$$

The first term in eqn. (11) encourages a softmax loss on the original training set S and the second term in eqn. (11) encourages a good prediction on the individual pseudo-negative class generated for the $k-th$ class. Note that we now only need to build a single CNN sharing $\mathbf{w}^{(0)}_t$ for all the K classes. We perform direct multi-class classification where the parameter setting is identical to a standard multi-class classification in CNN whereas an additional "not-real" class is created in (Salimans et al., 2016).

3.4. Analysis

We show the convergence of $p^-_t(\mathbf{x}) \overset{t=\infty}{\to} p^+(\mathbf{x})$, inspired by the proof from (Tu, 2007),

Theorem 1 *$KL[p^+(\mathbf{x})||p^-_{t+1}(x)] \leq KL[p^+(\mathbf{x})||p^-_t(\mathbf{x})]$ where KL denotes the Kullback-Leibler divergence, and $p(x|y=+1) \equiv p^+(x)$, under the assumption that classifier at $t+1$ improves over that at t.*

Proof:

$$p^-_t(\mathbf{x}) = \frac{1}{Z_t} \frac{q_t(y=+1|\mathbf{x})}{q_t(y=-1|\mathbf{x})} p^-_r(\mathbf{x}),$$

and

$$p^-_{t+1}(\mathbf{x}) = \frac{1}{Z_{t+1}} \frac{q_{t+1}(y=+1|\mathbf{x})}{q_{t+1}(y=-1|\mathbf{x})} p^-_r(\mathbf{x}).$$

$$p^-_{t+1}(\mathbf{x}) = \frac{1}{H_{t+1}} \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} p^-_t(\mathbf{x})$$

where $H_{t+1} = \int \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} p^-_t(\mathbf{x}) d\mathbf{x}$

$$\begin{aligned} & KL[p^+(\mathbf{x})||p^-_t(\mathbf{x})] - KL[p^+(\mathbf{x})||p^-_{t+1}(\mathbf{x})] \\ = & \int p^+(\mathbf{x}) \ln \left(\frac{1}{H_{t+1}} \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} p^-_t(\mathbf{x}) \right) d\mathbf{x} \\ - & \int p^+(\mathbf{x}) \ln[p^-_t(\mathbf{x})] dx \\ = & \int p^+(\mathbf{x}) \ln \frac{1}{H_{t+1}} d\mathbf{x} \\ + & \int p^+(\mathbf{x}) \ln \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} d\mathbf{x} \\ = & \ln \frac{1}{H_{t+1}} + \int p^+(\mathbf{x}) \ln \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} d\mathbf{x} \geq 0. \end{aligned}$$

Since $H_{t+1} = \int \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} p^-_t(\mathbf{x}) d\mathbf{x} \leq 1$ and $\int p^+(\mathbf{x}) \ln \frac{\exp\{\mathbf{w}^{(1)T}_{t+1} \phi(\mathbf{x}; \mathbf{w}^{(0)}_{t+1})\}}{\exp\{\mathbf{w}^{(1)T}_t \phi(\mathbf{x}; \mathbf{w}^{(0)}_t)\}} d\mathbf{x} \geq 0$. We assume that classifier at $t+1$ improves over that at t. This shows that $p^-_{t+1}(\mathbf{x})$ converges to $p(\mathbf{x}|y=+1)$ and the convergence rate depends on the classification error at each step.

Remark Given a training set $S_{train} = \{(\mathbf{x}_i, y_i), i = 1..n\}$ and test set $S_{test} = \{(\mathbf{x}_i, y_i), i = 1..b\}$ where $y_i \in \{-1, +1\}$ and prediction function $f(\mathbf{x}) \in \{-1, +1\}$. Typically one computes a training error $\epsilon_{train} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}(y_i \neq f(\mathbf{x}_i))$ and a test error $\epsilon_{test} = \frac{1}{b} \sum_{i=1}^{b} \mathbf{1}(y_i \neq f(\mathbf{x}_i))$. The difference between ϵ_{train} and ϵ_{test}, $\epsilon_{test} = \epsilon_{train} + \epsilon_{generalization}(f)$ is largely due to the difference of the data distributions in training and testing. We pay particular attention to the negative samples, which are in a space that is often much larger than the positive sample space. For the negative training samples, we have $y_i = -1$ and $\mathbf{x}_i \sim Q^-(\mathbf{x})$, where $Q^-(\mathbf{x})$ is a distribution on the given negative examples in the original training set. Our reclassification-by-synthesis algorithm (Algorithm 1) essentially constructs a mixture model $\tilde{p}(\mathbf{x}) \equiv \frac{1}{T} \sum_{t=0}^{T-1} p^-_t(\mathbf{x})$ by sequentially generating pseudo-negative samples to augment our training set. Our new distribution for augmented negative sample set thus becomes

$$Q^-_{new}(\mathbf{x}) \equiv \frac{n}{n+Tl} Q^-(\mathbf{x}) + \frac{Tl}{n+Tl} \tilde{p}(\mathbf{x}),$$

where $\tilde{p}(\mathbf{x})$ encodes pseudo-negative samples that are confusing and similar to (but are not) the true positives. Often $th_1 \leq Dist(p^+(\mathbf{x}), \tilde{p}(\mathbf{x})) \leq th_2$ where th_1 and th_2 are two thresholds bounding the distances. Our overall algorithm thus is capable of enhancing classification by self-generating confusing samples to improve the robustness.

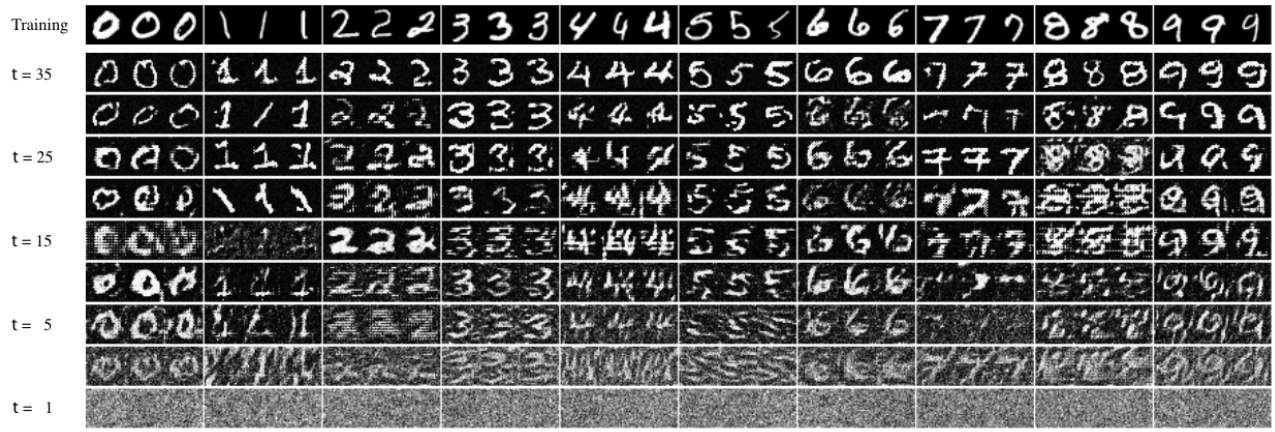

Figure 3. Synthesized pseudo-negatives for the MNIST dataset by our ICL classifier. The top row shows some training examples. As t increases, our classifier gradually synthesize pseudo-negative samples that become increasingly faithful to the training samples.

4. Experiments

We conduct experiments on three standard benchmark datasets, including MNIST, CIFAR-10 and SVHN. We use MNIST as a running example to illustrate our proposed framework using a shallow CNN and show competitive results using a state-of-the-art CNN classifier, ResNet (He et al., 2016a) on CIFAR-10 and SVHN.

In our experiments, for the reclassification step, we use the SGD optimizer with mini-batch size of 64 (MNIST) or 128 (CIFAR-10 and SVHN) and momentum equal to 0.9; for the synthesis step, we use the Adam optimizer (Kingma & Ba, 2014) with momentum term β_1 equal to 0.5. All results are obtained by averaging multiple rounds.

4.1. MNIST

Table 1. Test error on the MNIST dataset. Our ICL method outperforms the CNN baseline. A two-step approach combining GDL+CNN achieves a test error of 0.89% (one-vs-all) and that combining DCGAN+CNN achieves a test error 0.86% (one-vs-all). See description in text for the details about our two-step experiments for GDL+CNN and DCGAN+CNN.

Method	One-vs-all (%)	Softmax (%)
CNN (baseline)	0.91	0.89
ICL-noise (ours)	0.92	0.87
ICL (ours)	0.78	0.81

The MNIST (LeCun & Cortes, 1998) dataset consists of 28×28 grey images from 10 different classes (0 to 9) with $60,000$ training and $10,000$ test samples. We use a simple network, containing 4 convolutional layers, each having a 5×5 filter size with 64, 128, 256 and 512 channels, respectively. These convolutional layers have stride 2, and no pooling layers are used. LeakyReLU activations (Maas et al., 2013) are used after each convolutional layer. The last convolutional layer is flattened and fed into a sigmoid output (in the one-vs-all case). This network architecture is similar to the discriminator in DCGAN (Radford et al., 2015), but batch normalization layers are not used.

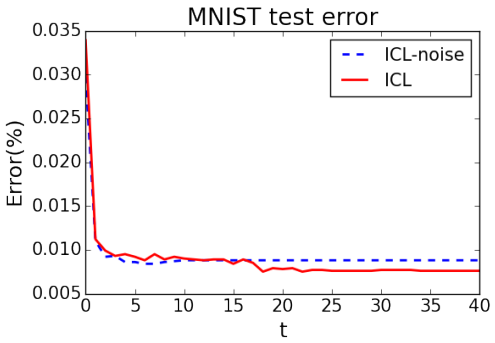

Figure 4. Test errors of ICL by the reclassification-by-synthesis algorithm on MNIST w.r.t. t during training.

In the synthesis step, we use the backpropagation sampling process as discussed in Section 3.2.2. We initialize the image with random noise, fix the parameters of the network, and run multiple steps of Adam with early stopping. We also implemented alternative sampling approaches but this early stopping has a good balance between effectiveness and efficiency (see discussions in Section 3.2.2). Each time we synthesize a fixed number (e.g. 200) of pseudo-negative samples. In the reclassification step, we run SGD (for 5 epochs) on the current training data S_e^t, including previously generated pseudo-negatives. Our initial learning rate is 0.01 and is decreased by a factor of 10 at $t = 20$.

We show some synthesized pseudo-negatives from the MNIST dataset in Figure 3. The samples in the top row are from the original training dataset. ICL gradually synthesizes pseudo-negatives, which are increasingly faithful to the original data. Since we treat these pseudo-negatives as negative samples in our reclassification step, we set a threshold on the total number of synthesis steps to use, as determined by cross-validation. For evaluation, we compare our ICL method with the baseline CNN models. In

the one-vs-all case, we train one network for each class individually, therefore we have 10 classifiers in total whereas in the softmax case, we only need to train a single network.

Comparison with GDL and GAN. GDL (Tu, 2007) focuses on unsupervised learning; GAN (Goodfellow et al., 2014a) and DCGAN (Radford et al., 2015) show results for unsupervised learning and semi-supervised classification. To apply GDL and GAN to the supervised classification setting, we design an experiment to perform a two-step implementation.

For GDL, we ran the GDL code (Tu, 2007) and obtained the pseudo-negative samples for each individual digit; the pseudo-negatives are then used as augmented negative samples to train individual one-vs-all CNN classifiers (using an identical CNN architecture to ICL for a fair comparison), which are combined to form a multi-class classifier in the end. To compare with DCGAN (Radford et al., 2015), we follow the same procedure: each generator trained by DCGAN (Radford et al., 2015) using tensorflow implementation (Kim, 2016) was used to generate positive samples, which are then augmented to the negative set to train the individual one-vs-all CNN classifiers (also using an identical CNN architecture to ICL), which are combined to create the overall multi-class classifier. GDL+CNN achieves a test error of 0.89% and DCGAN+CNN achieves a test error of 0.86% on the MNIST dataset, whereas ICL reports an error of 0.78% using the same CNN architecture. These results are averaged over multiple rounds. As the supervised learning task was not directly specified in DCGAN (Radford et al., 2015), some care is needed to design the optimal setting to utilize the generated samples from DCGAN in the two-step approach (we may not get the best setting in the experiments). GDL (Tu, 2007) can be made into a discriminative classifier by utilizing the given negative samples first but boosting (Freund & Schapire, 1997) with manually designed features was adopted which may not reproduce competitive results as CNN classifier does.

Nevertheless, the advantage of ICL being an integrated end-to-end supervised learning single-model framework can be observed.

Ablation study. We also experiment using random noise as the synthesized pseudo-negatives in an ablation study. Table 1 lists the test errors on the MNIST dataset over multiple trials. We observe our ICL method outperforms the CNN baseline and the ICL-noise method in both one-vs-all and softmax cases. The error plots of our algorithm are shown in Figure 4.

To better understand the effectiveness of our ICL method, we set up an experiment to see how it performs when varying the number of training examples. We vary the number of training examples across sets of size 500, 2000, 10000 and 60000. From Figure 5, we can observe that when we have fewer training examples, our ICL model is more effective since it has the internal synthesis capability to synthesize pseudo-negatives in order to aid the training process.

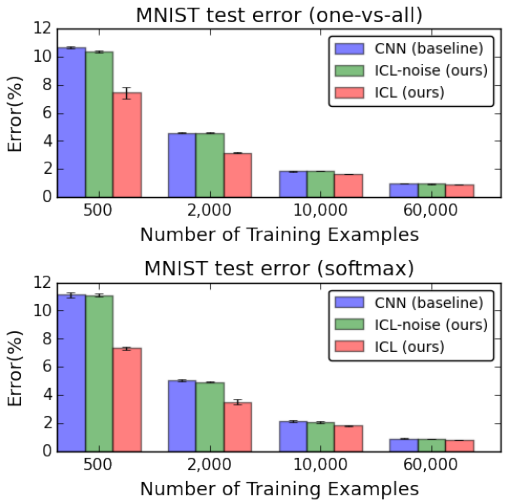

Figure 5. MNIST test error against the number of training examples (std dev. of the test error is also displayed). The effect of ICL is more clear when having fewer training examples.

Robustness to external adversarial examples. The synthesis strategy within ICL and its generative capability provides improvement that is intrinsically more robust to confusing and challenging examples. To validate this, we compare the baseline CNN with our ICL classifier on adversarial examples generated using the "fast gradient sign" method from (Goodfellow et al., 2014b) for validation purposes only. This "fast gradient sign" method (with $\epsilon = 0.25$) can cause a maxout network to misclassify 89.4% of adversarial examples generated from the MNIST test set (Goodfellow et al., 2014b). In our experiment, we set $\epsilon = 0.125$. Starting with $10,000$ MNIST test examples, we first determine those which are correctly classified by the baseline CNN in order to generate adversarial examples from them. We find that $6,876$ generated adversarial examples successfully fool the baseline CNN, however, only $4,661$ of these examples can fool our ICL classifier, which is a 32.2% reduction in error against adversarial examples. Note that the improvement is achieved without using any additional training data, nor knowing a prior about how these adversarial examples are generated by the specific "fast gradient sign method" (Goodfellow et al., 2014b). Figure 6 shows some examples. On the contrary, of the $5,989$ adversarial examples generated from the ICL classifier side that then fooled ICL using the same method, $5,542$ of them can still fool the baseline CNN classifier. This two-way experiment shows the significantly improved robustness of ICL over the baseline CNN.

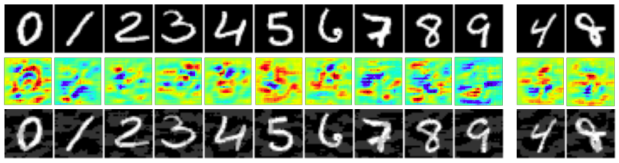

Figure 6. Validation on additional adversarial examples generated using "fast gradient sign method" (Goodfellow et al., 2014b) to show the improved robustness of ICL. The first row shows the original images, the second row displays the perturbations, and the third row shows the resulting adversarial examples. All these examples fooled the baseline CNN model and the last two columns show those that fooled ICL.

Table 2. Test error on the CIFAR-10 dataset. ResNet (He et al., 2016b) is adopted as our baseline [1]. In both one-vs-all and softmax cases, ICL shows an immediate improvement over the baseline ResNet.

Method	One-vs-all (%)	Softmax (%)
w/o Data Augmentation		
ResNet-32 (baseline)	13.44	12.38
ICL-noise (ours)	13.28	11.94
ICL (ours)	12.94	11.46
w/ Data Augmentation		
ResNet-32 (baseline)	6.70	7.06
ICL-noise (ours)	6.58	6.90
ICL (ours)	6.52	6.70

4.2. CIFAR-10

The CIFAR-10 (Krizhevsky, 2009) dataset consists of 32×32 color images. A collection of $60,000$ images is split into $50,000$ training and $10,000$ testing images. We adopt ResNet (He et al., 2016b) as our baseline model, which has been widely used. For data augmentation, we follow the standard procedure in (Lee et al., 2015; 2016; He et al., 2016b) by augmenting the dataset by zero-padding 4 pixels on each side; we also perform cropping and random flipping. In both one-vs-all and softmax cases, ICL outperforms the baseline method.

Our proposed ICL method is orthogonal to many existing approaches which use various improvements to the network structures in order to enhance the CNN performance.

4.3. SVHN

Table 3. Test error on the SVHN dataset. ICL outperforms the baseline ResNet (He et al., 2016b) model.

Method	Softmax (%)
ResNet-32 (baseline)	2.01
ICL-noise (ours)	1.99
ICL (ours)	1.93

The SVHN (Netzer et al., 2011) dataset consists of color images of house numbers collected by Google Street View. We use a training set that combines its training and extra data in our experiments and leave the test data as our test set. No data augmentation has been applied. Our ICL method shows improvement to the baseline ResNet.

[1] https://github.com/ppwwyyxx/tensorpack/tree/master/examples/ResNet

5. Conclusion

In this paper, we have proposed introspective classifier learning (ICL), which results in a CNN classifier empowered with generative capabilities. We developed a reclassification-by-synthesis training algorithm. A new loss function to seamlessly integrate ICL in a single multi-class CNN classifier is additionally developed. We tested the algorithm on standard benchmarks and report encouraging results.

6. Acknowledgement

This work is supported by NSF IIS-1618477 and a Northrop Grumman Contextual Robotics grant. We thank Saining Xie, Shuai Tang, and Sanjoy Dasgupta for helpful discussions.

References

Baldi, Pierre. Autoencoders, unsupervised learning, and deep architectures. *ICML unsupervised and transfer learning*, 27(37-50):1, 2012.

Beygelzimer, Alina, Dasgupta, Sanjoy, and Langford, John. Importance weighted active learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 49–56. ACM, 2009.

Breiman, Leo. Random Forests. *Machine Learning*, 2001.

Brock, Andrew, Lim, Theodore, Ritchie, JM, and Weston, Nick. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016.

Carreira-Perpinan, Miguel A and Hinton, Geoffrey. On contrastive divergence learning. In *AISTATS*, volume 10, pp. 33–40. Citeseer, 2005.

Chen, Tianqi, Fox, Emily B, and Guestrin, Carlos. Stochastic gradient hamiltonian monte carlo. In *ICML*, 2014.

Della Pietra, Stephen, Della Pietra, Vincent, and Lafferty, John. Inducing features of random fields. *IEEE transactions on pattern analysis and machine intelligence*, 19(4):380–393, 1997.

Dietterich, T. G., Lathrop, R. H., and Lozano, P. T. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89:31–71, 1997.

Duda, R. O., Hart, P. E., and Stork, D. G. *Pattern Classification*. 2nd edition, 2000.

Freund, Y. and Schapire, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. *J. of Comp. and Sys. Sci.*, 55(1), 1997.

Friedman, Jerome, Hastie, Trevor, and Tibshirani, Robert. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.

Gatys, Leon A, Ecker, Alexander S, and Bethge, Matthias. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*, 2015.

Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014a.

Goodfellow, Ian J, Shlens, Jonathon, and Szegedy, Christian. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014b.

Grenander, Ulf and Miller, Michael I. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 549–603, 1994.

Gutmann, Michael and Hyvärinen, Aapo. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, volume 1, pp. 6, 2010.

He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *CVPR*, 2016a.

He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pp. 630–645. Springer, 2016b.

Hinton, G. E., Osindero, S., and Teh, Y. W. A fast learning algorithm for deep belief nets. *Neural computation*, 18:1527–1554, 2006.

Hinton, Geoffrey E, Dayan, Peter, Frey, Brendan J, and Neal, Radford M. The" wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214):1158, 1995.

Jebara, Tony. Machine learning: discriminative and generative, 2012.

Kim, Taehoon. DCGAN-tensorflow. https://github.com/carpedm20/DCGAN-tensorflow, 2016.

Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Krizhevsky, Alex. Learning Multiple Layers of Features from Tiny Images. *CS Dept., U Toronto, Tech. Rep.*, 2009.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012.

Leake, David B. Introspective learning and reasoning. In *Encyclopedia of the Sciences of Learning*, pp. 1638–1640. Springer, 2012.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R.E., Hubbard, W., and Jackel, L.D. Backpropagation applied to handwritten zip code recognition. In *Neural Computation*, 1989.

LeCun, Yann and Cortes, Corinna. The MNIST database of handwritten digits, 1998.

Lee, Chen-Yu, Xie, Saining, Gallagher, Patrick, Zhang, Zhengyou, and Tu, Zhuowen. Deeply-supervised nets. In *AISTATS*, 2015.

Lee, Chen-Yu, Gallagher, Patrick W, and Tu, Zhuowen. Generalizing pooling functions in convolutional neural networks: Mixed, gated, and tree. In *AISTATS*, 2016.

Liang, Percy and Jordan, Michael I. An asymptotic analysis of generative, discriminative, and pseudolikelihood estimators. In *Proceedings of the 25th international conference on Machine learning*, pp. 584–591. ACM, 2008.

Liu, Jun S. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.

Maas, Andrew L, Hannun, Awni Y, and Ng, Andrew Y. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, 2013.

Mandt, Stephan, Hoffman, Matthew D, and Blei, David M. Stochastic gradient descent as approximate bayesian inference. *arXiv preprint arXiv:1704.04289*, 2017.

Mooney, Christopher Z, Duval, Robert D, and Duvall, Robert. *Bootstrapping: A nonparametric approach to statistical inference*. Number 94-95. Sage, 1993.

Mordvintsev, Alexander, Olah, Christopher, and Tyka, Mike. Deepdream - a code example for visualizing neural networks. *Google Research*, 2015.

Netzer, Yuval, Wang, Tao, Coates, Adam, Bissacco, Alessandro, Wu, Bo, and Ng, Andrew Y. Reading Digits in Natural Images with Unsupervised Feature Learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.

Quinlan, J.R. Improved use of continuous attributes in c4.5. *J. of Art. Intell. Res.*, 4, 1996.

Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

Salimans, Tim, Goodfellow, Ian, Zaremba, Wojciech, Cheung, Vicki, Radford, Alec, and Chen, Xi. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016.

Settles, Burr. Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11, 2010.

Sutton, Richard S and Barto, Andrew G. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.

Tolstikhin, Ilya, Gelly, Sylvain, Bousquet, Olivier, Simon-Gabriel, Carl-Johann, and Schölkopf, Bernhard. Adagan: Boosting generative models. *arXiv preprint arXiv:1701.02386*, 2017.

Tu, Zhuowen. Learning generative models via discriminative approaches. In *CVPR*, 2007.

Tu, Zhuowen, Narr, Katherine L, Dollár, Piotr, Dinov, Ivo, Thompson, Paul M, and Toga, Arthur W. Brain anatomical structure segmentation by hybrid discriminative/generative models. *Medical Imaging, IEEE Transactions on*, 27(4):495–508, 2008.

Vapnik, Vladimir N. *The nature of statistical learning theory*. Springer-Verlag New York, Inc., 1995.

Welling, Max and Teh, Yee W. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011.

Welling, Max, Zemel, Richard S, and Hinton, Geoffrey E. Self supervised boosting. In *Advances in neural information processing systems*, pp. 665–672, 2002.

Wu, Ying Nian, Zhu, Song Chun, and Liu, Xiuwen. Equivalence of julesz ensembles and frame models. *International Journal of Computer Vision*, 38(3), 2000.

Xie, Jianwen, Lu, Yang, Zhu, Song-Chun, and Wu, Ying Nian. Cooperative training of descriptor and generator networks. *arXiv preprint arXiv:1609.09408*, 2016a.

Xie, Jianwen, Lu, Yang, Zhu, Song-Chun, and Wu, Ying Nian. A theory of generative convnet. *arXiv preprint arXiv:1602.03264*, 2016b.

Zhao, Junbo, Mathieu, Michael, and LeCun, Yann. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.

Zhu, Song Chun, Wu, Ying Nian, and Mumford, David. Minimax entropy principle and its application to texture modeling. *Neural Computation*, 9(8):1627–1660, 1997.

Zhu, Xiaojin. Semi-supervised learning literature survey. *Technical Report 1530*, 2005.